

# A Literature Survey on Empirical Evidence in Software Engineering

Ekaterina Pek

Software Languages Team, ADAPT Lab  
Universität Koblenz-Landau, Germany

Ralf Lämmel

Software Languages Team  
Universität Koblenz-Landau, Germany

**Abstract—Context:** Software Engineering research makes use of collections of software artifacts (corpora) to derive empirical evidence from. **Goal:** To improve quality and reproducibility of research, we need to understand the characteristics of used corpora. **Method:** For that, we perform a literature survey using grounded theory. We analyze the latest proceedings of seven relevant conferences. **Results:** While almost all papers use corpora of some kind with the common case of collections of source code of open-source Java projects, there are no frequently used projects or corpora across all the papers. For some conferences we can detect recurrences. We discover several forms of requirements and applied tunings for corpora which indicate more specific needs of research efforts. **Conclusion:** Our survey feeds into a quantitative basis for discussing the current state of empirical research in software engineering, thereby enabling ultimately improvement of research quality specifically in terms of use (and reuse) of empirical evidence.

## I. INTRODUCTION

This is a survey on software engineering research with focus on the use of collections of software artifacts (corpora) to derive empirical evidence from. Such focus on corpora was triggered by our own research on specifically software reverse/re-engineering and program comprehension, e.g., studies on API or language usage [1], [2], [3], [4]—with the common use of corpora for validation in the broader sense. The survey applies to conferences that fit with this context. One can observe a diversity of involved methodologies and characteristics of the collections of empirical evidence as they are leveraged in SE research. Thus, we embarked on the present literature survey with the following central research questions:

- I How often do Software Engineering papers use *corpora*—collections of empirical evidence?
- II What is the nature and characteristics of the used corpora?
- III Does common contents occur in the used corpora?

For this, we collected and analyzed the latest proceedings<sup>1</sup> of the following conferences: European Conference on Software Maintenance and Reengineering (CSMR), International Symposium on Empirical Software Engineering and Measurement (ESEM), International Conference on Program Comprehension (ICPC), International Conference on Software Maintenance (ICSM), Working Conference on Mining Software Repositories (MSR), Working Conference on Source Code Analysis and Manipulation (SCAM), Working Conference on

Reverse Engineering (WCRE). We choose these conferences because i) they cover software engineering topics that, based on our experience, we expect to make use of empirical evidence; ii) they cover ground related to our expertise and research focus on software reverse/re-engineering and program comprehension with ESEM as notable addition for broader coverage of empirical software engineering research; iii) the conferences are of comparable size. In our survey, we use only long papers. We choose to analyze only conference proceedings, because while journal articles may adhere to the best practices, conference proceedings arguably contain the most common practices of research in the community—and we are interested in the latter.

SE research has been surveyed before; see Table I for a summary. The cited surveys focus on specific forms or characteristics of SE research to be analyzed with a predefined schema. For instance, Kitchenham et al. surveyed SE journals and conferences to find out adoption rate of systematic literature reviews [8]. Similarly, Sjøberg et al. sought to find and analyze existing controlled experiments in SE research [6]. By contrast, we (first to our knowledge) seek to discover whatever empirical evidence is used to facilitate SE research and we allow our coding schema to emerge from the data. We follow the idea of Grounded Theory (GT) as understood by Glaser [9] (on the difference between Straussian and Glaserian versions see [10]).

The paper is organized as follows: §II describes the methodology underlying this literature survey. §III presents the results of the survey. §IV discusses related surveys. §V identifies threats to validity. §VI concludes the paper.

## II. METHODOLOGY

Empirical research is usually perceived as taking one of the forms: controlled and quasi-experiments, exploratory and confirmatory case studies, survey, ethnography, and action research [11], [12]. In a broader sense, empirical research also includes any research based on collected evidence—quoting from [11]: “Empirical research seeks to explore, describe, predict, and explain natural, social, or cognitive phenomena *by using evidence based on observation or experience*. It involves obtaining and interpreting evidence by, e.g., experimentation, systematic observation, interviews or surveys, or *by the careful examination of documents or artifacts* [emphasis added].”

Since Software Engineering is a practical area of Computer Science, it is logical to expect that most of the SE research is evidence-based, i.e., empirical de facto and in the

<sup>1</sup>As they are available from the DBLP bibliography service, <http://dblp.uni-trier.de/>

TABLE I. LITERATURE SURVEYS ON SOFTWARE ENGINEERING RESEARCH

Name	Ref	Year	# used		Period	# papers			Focus	Coding schema
			j	c		total	sel.	rel.		
Glass et al.	[5]	2002	6	0	1995–1999	—	369	369	Characteristics of SE research	Topics, research approaches and methods, theoretical basis, level of analysis
Sjøberg et al.	[6]	2005	9	3	1993–2002	5453	103	103	Controlled experiments	Extent, topic, subjects, task and environment, replication, internal and external validity
Zannier et al.	[7]	2006	0	1	1975–2005	1227	63	44	Empirical evaluation: quantity and soundness	Study type, sampling type, target and used population, evaluation type, proper use of analysis, usage of hypotheses
Kitchenham et al.	[8]	2009	10	3	2004–2007	2506	33	19	Systematic reviews	Inclusion and exclusion criteria, coverage, quality/validity assessment, description of the basic data
<b>Our study</b>		<b>2013</b>	<b>0</b>	<b>7</b>	<b>2011/2012</b>	<b>227</b>	<b>175</b>	<b>175</b>	<b>Empirical evidence</b>	<b>Emerged classification</b>

Legend: *j* and *c* stand for *journals* and *conferences*; *sel.* and *rel.* stand for *selected* and *relevant*.

TABLE II. CONFERENCES USED IN THE SURVEY

Year	Conference	# papers	
		total	long
2012	CSMR	30	30
2012	ESEM	43	24
2012	ICPC	23	21
2011	ICSM	36	36
2012	MSR	29	18
2011	SCAM	19	19
2011	WCRE	47	27
Total		227	175

present study, we submit to substantiate this expectation. We believe that a bottom-up approach of observing what exists and discovering methodology as well as definitions of forms of research complements the prominent top-down approach, when a methodology is derived from theoretical considerations or by borrowing from other sciences (medicine, sociology, psychology).

This survey is particularly concerned with (collections of) empirical evidence. Thus, the following questions guide the research:

- I How often do Software Engineering papers use *corpora*—collections of empirical evidence?
- II What is the nature and characteristics of the used corpora?
- III Does common contents occur in the used corpora?

For that, we collected the papers from the latest edition of seven SE conferences: CSMR, ESEM, ICPC, ICSM, MSR, SCAM, and WCRE (see Table II for details). We used DBLP pages of conferences to identify long papers and downloaded them from digital libraries.

We then proceeded to read the papers to perform coding. From a previously done, smaller and more specific literature survey [13] and a pilot study for the present survey, we had some basic understanding of the parts of the scheme to emerge. During the first pass of coding, we started with the empty scheme and completed it eventually to arrive at the current scheme, as described below. During the second pass, we compared profiles of coded papers against the latest version of the scheme, we went through the papers again and filled in the missing details.

While we were interested primarily in characteristics of used empirical evidence (specifically corpora), we also extracted additional information about research reported in the papers: used tools, signs of rigorousness/quality, etc. We put the collected information in several groups:

1) *Corpora*: We captured what was used as study objects (e.g., projects), what are their characteristics (e.g., language, open- vs. closed-source, code form), what are the requirements to the study objects, do they come from a specific source (e.g., established dataset or online repository), were they observed over a time (e.g., versions or revisions), what is the nature of preparation of the corpus.

2) *Forms of empirical research*: During coding, several structural forms evolved that we used for capturing information conveyed in papers: experiments, questionnaires, literature surveys, and comparisons. Some relationships between forms and corpora usage also emerged.

3) *Self-classification*: For each paper we captured what words authors use to describe their effort: e.g., case study, experiment.

4) *Tools*: We collected mentions of existing tools (e.g., Eclipse, R, Weka) that were used as well as of introduced tools that were presented in the papers. (In many cases, these tools are used to analyze or to otherwise process corpora.)

5) *Structural signs of rigorousness/quality*: We paid attention to the following aspects of the study presentation: Do authors use research questions? Null hypotheses? Is there a section on definitions and terms? Is validation mentioned? Is there a “Threats to validity” section? Are threats addressed in any structured way?

6) *Reproducibility*: We tried to understand in each case, if a study can be reproduced. (Obviously, the use of corpora affects the definition of reproducibility.) We paid attention to the following signs: Are all details provided for a possible study replication (i.e., versions of used projects, time periods, etc.)? Do authors provide any material used in the paper, e.g., on a supplementary website? Altogether, would it be possible to reproduce the study?

7) *Assessment*: Finally, we characterized the process of coding: how easy it was to extract information and how confident we are in the result.

We did the pilot survey in September-October 2012. After that, we adjusted our methodology (e.g., instead of filtering papers based on their abstract, we decided to survey all the papers) and proceeded to perform the current study in November 2012-January 2013. We used Python and Bash scripts, Google Refine tool<sup>2</sup>, and R project<sup>3</sup> to process the data. We provide online the list of the papers and results of coding<sup>4</sup>.

### III. RESULTS

In this section, we present the results of our study. We group them similarly to the description provided in Section II: details about detected corpora, emerged forms of empirical research, used or introduced tools, signs of rigorousness/quality of research, reproducibility of the studies, and, finally, assessment of our effort. When we use the phrase “*on the average*”, we imply the median of the appropriate distribution.

Next to the numbers, we provide framed highlights.

We use formula “X out of Y papers” to provide feeling for the numbers. E.g., “one out of three papers” means that in every three surveyed papers there is one that has the discussed characteristic.

We also provide conference-wise percentage of found characteristics. The table below illustrates the format on an artificial example: conferences are listed from left to right as the percentage increases. Percentage is always given relative to the total number of the long papers in the conference. Where appropriate, below the percentage appear names of the most popular projects, requirements, tunings within the conferences. When more than one name is given, each of them appear with the specified frequency.

Artificial example						
CSMR	ESEM	ICPC	ICSM	MSR	SCAM	WCRE
1 %	2 %	3 %	4 %	5 %	6 %	7 %

#### A. Corpora

1) *Usage*: We marked a paper as containing a corpus when the paper mentioned a collection of software artifacts used for deriving empirical evidence. Altogether, we have found 198 corpora used in 165 papers out of 175 surveyed papers.

In 28 cases, we decided that a paper contains more than one corpus. We did so consistently, when we met at least two of the following motivations mentioned in the paper when describing the purpose of collected empirical evidence: for benchmark or oracle (6 corpora), for training (6 corpora), for evaluation (5 corpora), for investigation (5 corpora), for testing (4 corpora), for investigating quality like accuracy or scalability (4 corpora).

We have found that 168 corpora (used in 145 papers), consist of projects (systems, software); in other cases, corpora consist of another kind of study object: image, trace, feature,

web log, etc. Till the end of the current subsection (III-A), we restrict ourselves to the corpora consisting of projects and call them project-based corpora.

Almost all papers use a corpus of some sort. One out of six papers has more than one corpus. Most of the corpora consist of projects.

Project-based corpora usage						
ESEM	ICPC	WCRE	SCAM	CSMR	MSR	ICSM
58 %	81 %	81 %	84 %	87 %	89 %	94 %

2) *Contents*: We identified the following common characteristics of project-based corpora.

**Size.** Half of the corpora, 99 cases, have three or less projects (of them, 45 corpora consist of only one project). There are 24 corpora that contain more than 10 projects. We detected large corpora (with more than 100 projects) in 8 papers—one of them introducing an established dataset itself.

**Languages.** Most of the corpora are monolingual (147 cases); most of the remaining ones are bilingual (19 cases). As for the software language, 106 corpora contain projects written in Java, while C-like languages are used in 50 corpora (in C-like languages we include C, C++, C#).

**Code form.** In 125 cases, corpora consist of source code; in 15 cases—of binaries. In the rest of the cases, code of the projects is not used, something else is in focus (developers, requirements, etc.)

**Access.** In 128 cases, corpora consist only of open-source projects; in 12 cases, corpora consist only of projects not available publicly (e.g., industrial software); in 9 cases, corpora are self-written. The remaining cases mix access forms.

**Projects.** We collected names of the used projects as they are provided by the papers (modulo merging of names like Vuze/Azureus<sup>5</sup>).

TABLE III. USED PROJECTS

Project	# corp
JHotDraw	15
JEdit	12
Ant	11
ArgoUML	11
Eclipse	11
Firefox	10
Vuze/Azureus	8
Linux kernel	6
Lucene	6
Mozilla	6
Hibernate	5

Table III lists projects frequently used in the corpora. Eclipse is a complex project, and some corpora make use of its sub-parts, considering them as projects on their own (e.g., JDT Search, PDE Build)—counting such cases, there are altogether 22 papers making use of Eclipse.

**Units.** We captured when some unit related to the project was in the focus of the study: a bug report or a UML class diagram—namely, we would capture the fact when such unit was used to give quantitative information (e.g., in a table presenting number of bug reports in the project under investigation). The most popular units turned out to be bug reports, they are used in 21 cor-

<sup>2</sup><http://code.google.com/p/google-refine/>

<sup>3</sup><http://www.r-project.org/>

<sup>4</sup><http://softlang.uni-koblenz.de/empsurvey>

<sup>5</sup>The project changed its name in 2008.

TABLE IV. ONLINE REPOSITORIES AND ESTABLISHED DATASETS

Repository	# papers	Ref Dataset	# papers
SourceForge <sup>1</sup>	6	[14] SIR	3
Apache.org <sup>2</sup>	3	[15] MSR challenge	2
GitHub <sup>3</sup>	3	[16] P-MARt	2
Android Market <sup>4</sup>	2	[17] PROMISE	2
CodePlex <sup>5</sup>	2	[18] Qualitas	2

<sup>1</sup> <http://sourceforge.net/><sup>2</sup> <http://projects.apache.org/><sup>3</sup> <https://github.com/><sup>4</sup> Now known as Google Play, <https://play.google.com/store><sup>5</sup> <http://www.codeplex.com/>

pora; defects (faults, failures) are used in 16 corpora; tests—in 10; traces—in 5.

An average project-based corpus consists of source code of three open-source projects, written in Java. Eclipse or its sub-parts is used in one out of eight papers using project-based corpora. The projects used in at least five papers are JHotDraw, JEdit, Ant, ArgoUML, Firefox, Vuze/Azerus, Linux kernel, Lucene, Mozilla, and Hibernate. Within the corpus, bug reports, defects, tests, and traces can be in the focus of the study.

#### Popular projects

WCRE	SCAM	ESEM	ICSM	CSMR	ICPC	MSR
11 %	11 %	13 %	14 %	23 %	24 %	28 %
Eclipse	Lynx	Eclipse	ArgoUML	Eclipse	JEdit	Firefox
JEdit	Minix					Eclipse

**Sources.** When papers clearly state the source of their corpora, we collected such information.

Online repositories used in more than one paper are listed in Table IV. The rest of detected online repositories are used in only one paper each: BlackBerry App World<sup>6</sup>, Google Code<sup>7</sup>, Launchpad<sup>8</sup>, ShareJar<sup>9</sup>.

Established datasets used in more than one paper are listed in Table IV. Some of the other datasets that used only in one paper each: Bug prediction dataset [19], CHICKEN Scheme benchmarks<sup>10</sup>, CoCoMe<sup>11</sup>, DaCapo [20], FLOSSMetrics<sup>12</sup>, iBUGS<sup>13</sup>, SMG2000 benchmark<sup>14</sup>, SourcererDB [21], TEFSE challenge<sup>15</sup>. Table V summarizes the most popular types of sources and their distribution across conferences.

One out of four project-based corpora uses an established dataset, previous work, or online repository as a source of the projects. There is no common frequently

used dataset or repository. Only SourceForge shows moderately frequent usage.

#### Usage of corpora sources

ESEM	SCAM	ICPC	ICSM	WCRE	CSMR	MSR
13 %	16 %	24 %	25 %	26 %	30 %	39 %

**3) Evolution:** We encountered 52 papers that use evolution of the projects in their research, meaning that they operate on several versions, releases, etc. To describe the evolution measure, the following terms were used: “version” (21 times), “revision” (11), “commit” (10), “release” (11).

On the average, papers mentioning commits use 3,292 commits; papers with revisions—18,870 revisions; with versions—10 versions; with releases—10 releases.

There are 46 papers that mention a time span of their study. In 36 cases, the unit of the time span is a year and on the average such papers are concerned with a 8-year span.

We found 23 papers to mention what version control system was involved in the study. CVS is mentioned 11 times, SVN—11 times, Git and Mercurial—4 and 2 times respectively.

One out of three papers with project-based corpora uses evolution aspect in its research. In half of the cases, large-scale evolution is involved: several thousands commits/revisions or ten versions/releases of projects—often spanning several years of a project’s lifetime.

#### Evolution usage

ICPC	SCAM	ESEM	CSMR	ICSM	WCRE	MSR
14 %	16 %	21 %	33 %	33 %	33 %	56 %

**4) Requirements:** We collected requirements to the corpora: explicit as well as implicit. For instance, an implicit requirement for a bug tracking system is inferred if the paper uses bug reports of the projects under investigation. The most popular direction of requirements is the presence of some ‘ecosystem’ (found in 37 papers): existence of bug tracking systems, mailing lists, documentation (e.g., user manuals). Another popular requirement, found in 25 papers, has to do with the size of the projects: small, sufficient, large, or of particular size (as specific as “medium of the sizes of the ten most popular Sourceforge projects”), or the need of diversity of sizes. In 23 papers, it was stated that the used projects were chosen because they were used in previous work (of the same or other authors). Language-related requirement was present in 22 papers for a specific language or for the diversity of languages in a corpus. In 14 papers, the choice of projects was attributed to either diversity of application domains or to a specific domain. Some aspect of the used projects was mentioned as essential in 14 papers: active or wide-spread usage, popularity, well-known and established software. Other popular requirements include presence of development history (15 papers), dependencies (11 papers), or tests (10 papers).

<sup>6</sup><http://appworld.blackberry.com/webstore><sup>7</sup><http://code.google.com/><sup>8</sup><https://launchpad.net/><sup>9</sup><http://www.sharejar.com/><sup>10</sup><https://github.com/mario-goulart/chicken-benchmarks><sup>11</sup><http://agrausch.informatik.uni-kl.de/CoCoMe><sup>12</sup><http://libresoft.es/research/projects/flossmetrics><sup>13</sup><http://www.st.cs.uni-saarland.de/ibugs/><sup>14</sup>[https://asc.llnl.gov/computing\\_resources/purple/archive/benchmarks/smg/](https://asc.llnl.gov/computing_resources/purple/archive/benchmarks/smg/)<sup>15</sup><http://www.cs.wm.edu/semeru/tefse2011/Challenge.htm>

TABLE V. SOURCES OF CORPORA

Type	# papers							
	Total	CSMR	ESEM	ICPC	ICSM	MSR	SCAM	WCRE
Established dataset	20	5	0	2	6	5	0	2
Previous work	13	2	3	2	1	1	2	2
Online repository	12	3	0	1	2	2	1	3
Total	43	9	3	5	9	7	3	7
Percentage	25	30	13	24	25	39	16	26

One out of five papers requires the projects of its corpus to have an ecosystem: a bug tracker, or a mailing list, or some kind of documentation. Other requirements focus on the size and language of the projects, application domain, development history, etc.

Popular requirements

SCAM	MSR	ICSM	ICPC	ESEM	WCRE	CSMR
11 %	28 %	14 %	24 %	13 %	11 %	23 %
domain	ecosys	size	size	ecosys	ecosys	ecosys
lang	p.work					
size						

5) *Tuning*: We captured what kind of action is applied to a corpus during research. In 20 papers, sources or binaries were modified by instrumentation, faults/clones injection, adjusting identifiers, etc. In 15 papers, tests needed to be run against the corpus either to verify made modifications or to collect the data. In 10 papers, the corpora had to be executed in order to perform the needed analysis or to collect data. In 6 papers, some filtering of the contents of the corpus was needed to, e.g., identify main source code/main part of the project.

We have detected few common actions applied to corpora during research: source code/binaries modification; execution of the tests on the corpus or of the corpus itself; filtering of the corpus contents. Altogether, one out of four papers contains signs of one of these actions.

Popular actions

ESEM	MSR	SCAM	ICPC	ICSM	WCRE	CSMR
8 %	11 %	11 %	14 %	19 %	19 %	20 %
tests	run	modif.	modif.	tests	modif.	modif.

We captured manual effort that went into creation of a corpus, e.g., when a paper mentions setting up environments and providing needed libraries in order to execute the corpus. For that, we graded each corpus on the following scale. *None*: no manual effort mentioned (120 corpora); *some*: some manual effort mentioned, e.g., manual detection of design patterns in source code (33 corpora); and *all* means that corpus is self-written (10 corpora).

One out of four project-based corpora requires some manual effort.

Manual effort

CSMR	ESEM	ICSM	SCAM	ICPC	MSR	WCRE
13 %	13 %	22 %	26 %	33 %	33 %	37 %

### B. Self-classification

TABLE VI. SELF CLASSIFICATION

Type	#
case study	48
experiment	44
empirical study	22
evaluation	14
exploratory study	6
...	...

We collected explicit self-classifications from the papers; from the sentences like “we have conducted a case study” we would conclude that the current paper is a case study. Some of the self-classifications were very detailed and precise, e.g., “a pre/post-test quasi experiment”, in such cases we reduced the type to a simpler version, e.g., an experiment. We would also count terms like “experimental assessment” or “experimental study” towards the experiment type. As seen from Table VI, most often authors use terms such as “case study” and “experiment” to describe their research. In some cases, papers contain more than one self-classification (24 cases). In 36 papers, we could not detect any self-classification.

Four out of five papers provide self-classification, but it might be vague. The most popular term, ‘case study,’ may be misused. Cf., “There is much confusion in the SE literature over what constitutes a case study. The term is often used to mean a worked example. As an empirical method, a case study is something very different.” [12]. Cf., “... our sample indicated a large misuse of the term case study.” [7]

Self-classification

SCAM	MSR	CSMR	WCRE	ICPC	ESEM	ICSM
37 %	61 %	80 %	81 %	86 %	88 %	94 %

### C. Emerged forms

Independently of the self-classification of the papers, we noted structural characteristics of research performed in the papers. We did not use any theoretical definition for what to consider a questionnaire or an experiment. The developed definitions are structural, composed of the characteristics that

emerged from the papers, as they were discussed and structurally supported by the authors.

1) *Experiment*: We have identified 22 experiments in 19 papers. Except for two, they all involve human subjects. On the average, an experiment has 16 participants. The maximum number of participants is 128, the minimum is 2, first and third quartiles are 5 and 34 respectively. In 21 cases, an experiment uses a corpus (in 17 cases, a project-based one); 20 questionnaires are used in 10 experiments.

In two-thirds of the experiments, participants come from one population, the remaining experiments draw participants from two or three populations. The most common source of participants is students; sometimes distinguished by their level—graduate, undergraduate, Bachelor, Master, and PhD students. In one-third of the cases, professionals are involved (full-time developers, experts, industry practitioners, etc.). In half of the cases, participants form the only group in the experiment. When there is more than one group (usually, two—with a couple of exceptions of 4 and 5 groups), the group is representing a treatment (a task), or an experience level, or a gender. On the average, an experiment has 4 tasks and lasts for an hour (with a few exceptions when an experiment takes several weeks or even a month).

In 6 cases, it is mentioned that an experiment had a pilot study. In 6 cases, it is mentioned that participants of the experiment were offered compensation: monetary or another kind of incentive (e.g., a box of candy).

The main requirement for the participants is their experience: basic knowledge of used technology, or language, or IDE. As for the tasks, they are expected to be of a certain size (e.g., a method body to fit on one page), or of certain contents (e.g., contain “if” statements). The usual requirement for an experiment also is either that the tested tool or used code is unfamiliar to the participants, or on the contrary that the background is familiar (e.g., well-known design patterns).

One out of ten papers contains an experiment. The majority of the experiments use project-based corpora; experiments often use questionnaires, usually two per experiment. An average experiment involves 16 students, often in two groups (by the received treatment or experience level); it consists of four tasks and lasts for an hour. One out of four experiments suggests some compensation to its participants; one out of four experiments is preceded by a pilot study.

ICPC and ESEM are the main source of experiments involving professionals.

Experiments						
MSR	SCAM	CSMR	WCRE	ICSM	ESEM	ICPC
0 %	0 %	3 %	7 %	8 %	21 %	38 %

2) *Questionnaire*: Altogether, we have found 36 questionnaires in 24 papers. As mentioned, 20 questionnaires are used in experiments—to distinguish, we will refer to them as experiment-related and the other 16 we will qualify as experiment-unrelated.

Size-wise, there is no particular difference between experiment-related and -unrelated questionnaires. On the average, both have 20 questions grouped in one section. In 6 cases, an experiment-unrelated questionnaire has a corpus.

While experiment-related questionnaires have the same participants as the experiments they relate to (i.e., involve mostly students), experiment-unrelated questionnaires involve professionals (testers, managers, experts, consultants, software engineers) as participants in two-thirds of the cases. On the average, an experiment-unrelated questionnaire has 12 participants. When it was possible (6 cases), we calculated how many participants took part in the experiment-unrelated questionnaire compared to the initial number of questioned people. On the average, 19 % take part in the end, in the worst case the ratio can be as low as 5 %.

While experiment-related questionnaires have the same requirements regarding the participants as the experiments they relate to, experiment-unrelated questionnaires have requirements concerned with the participants’ experience (e.g., Java experience) or expertise (specific area of experience such as clone detection or web development).

When related to experiments, questionnaires are often performed before (referred to as “pretest” in 6 cases) and after the experiment (referred to as “posttest” in 9 cases).

In 5 cases, an experiment-unrelated questionnaire was preceded by a pilot study.

More than half of the detected questionnaires are used in experiments—often as pretest and posttest questionnaires. The other half, experiment-unrelated questionnaires, are found in one out of twelve papers. Size-wise, on the average there is no difference between experiment-related and -unrelated questionnaires. Experiment-unrelated questionnaires usually involve professionals as participants—in contrast to experiment-related questionnaires that mostly use students. Typical requirements for participants in experiment-unrelated questionnaires have to do with experience or expertise. One out of three experiment-unrelated questionnaires are preceded by a pilot study.

Experiment-unrelated questionnaires						
MSR	CSMR	SCAM	WCRE	ICSM	ICPC	ESEM
0 %	3 %	5 %	7 %	8 %	19 %	25 %

3) *Literature survey*: We have found 6 literature surveys in 5 papers. Except for one, they provide extensive details on how the survey was conducted. In particular, the used methodology is clearly stated: four times it is said to be a “systematic literature review” and once a “quasi systematic literature review”. In three cases, the systematic literature review was done following guidelines by Kitchenham [22].

The papers are initially collected either by searching digital libraries or from the proceedings of specific conferences and journals. Among used digital libraries are EI Compendex, Google Scholar, ISI, and Scopus—the latter was used in two papers. As for the conferences and journals, there is

no intersection between the lists of names—except for ICSE, which was used in two papers.

On the average, a literature survey starts with 2161 papers, its final set contains 35 papers, meaning that on the average only 1.6% papers are taken into account in the end. The percentage can be as high as 39% and as low as 0%.

Requirements for papers to be included into the survey are usually related to the scope of the investigated research. Other requirements are concerned with the paper itself: available online, written in English, a long paper, with empirical validation.

After all the papers are collected, they are filtered based on the titles and abstracts, which are examined manually by the researchers (in one case, also conclusions were taken into account; in another case, full text of the papers was searched for keywords). Then the full text of each paper is read and the final decision is made as to whether to consider the paper relevant.

Literature surveys are quite rare: only one out of 35 papers contains it. On the average, a literature survey starts with few thousand papers to be filtered down to few dozens papers that will be analyzed. Usually, the first round of filtering is based on the title and abstract, then the full text of the papers is considered. There is not enough information to conclude about frequently used digital libraries or conferences/journals. Half of the surveys were following guidelines of systematic literature reviews by Kitchenham [22].

Literature surveys						
ICSM	MSR	SCAM	WCRE	CSMR	ICPC	ESEM
0%	0%	0%	0%	3%	5%	13%

4) *Comparisons*: During coding, we noticed the recurring motif of comparisons in the papers. While we did not assess the scope nor the goal, we have coded the basic information: what is the nature of the subjects being compared (tools, techniques), how many subjects are compared, and is one of them introduced in the paper.

We have found comparisons in 56 papers. Almost all of them (except for 5 papers), use project-based corpora. Half of the time, a comparison is made for the technique, approach, or tool that was introduced in the study—with the apparent reason to evaluate the proposed technique, approach, or tool. On the average, such evaluation involves one other technique, approach, or tool. In the other cases, compared were: metrics, tools, algorithms, designs, etc. For such comparisons, on the average, the group of compared entities was of size 3.

One out of three papers compares tools, techniques, approaches, metrics, etc.—half of the time, to evaluate what was introduced in the study. On the average, such evaluation involves one other entity. In the other half of the cases, the average number of compared entities is 3.

TABLE VII. EXISTING TOOLS USED IN THE PAPERS

Tool	# papers	Tool	# papers
Eclipse <sup>1</sup>	25	MALLET <sup>7</sup>	4
R project <sup>2</sup>	16	ChangeDistiller <sup>8</sup>	3
CCFinder <sup>3</sup>	6	CodeSurfer <sup>9</sup>	3
Understand <sup>4</sup>	6	Evolizer <sup>10</sup>	3
Weka <sup>5</sup>	6	RapidMiner <sup>11</sup>	3
ConQAT <sup>6</sup>	4	RECODER <sup>12</sup>	3

<sup>1</sup> <http://eclipse.org>

<sup>2</sup> <http://www.r-project.org/>

<sup>3</sup> <http://www.ccfinder.net/>

<sup>4</sup> <http://www.scitools.com/>

<sup>5</sup> <http://www.cs.waikato.ac.nz/ml/weka/>

<sup>6</sup> <https://www.conqat.org/>

<sup>7</sup> <http://mallet.cs.umass.edu/>

<sup>8</sup> <http://www.ifi.uzh.ch/seal/research/tools/changeDistiller.htm>

<sup>9</sup> <http://www.grammatech.com/products/codesurfer/overview.html>

<sup>10</sup> <http://www.ifi.uzh.ch/seal/research/tools/evolizer.html>

<sup>11</sup> <http://rapid-i.com/content/view/181/190/>

<sup>12</sup> <http://sourceforge.net/projects/recoder/>

#### Comparisons

ICPC	MSR	SCAM	WCRE	ESEM	ICSM	CSMR
19%	22%	26%	30%	33%	36%	47%

#### D. Tools

We have found 46 papers to introduce a tool (where we were able to capture this fact only if the name of the tool was mentioned or it was clearly stated that “a prototype” is implemented). In 46 more papers, we detected that additional, helper tooling for the current purpose of the study is implemented (parsers, analyzers, and so on).

When names of existing tools were explicitly mentioned to be used, we collected the names. We have found that in 126 cases, a paper makes use of existing tools. On the average, a paper uses 2 tools; the captured maximum is 6. The frequently used tools are listed in Table VII. We counted towards Eclipse usage also cases when a paper used an existing tool that we know to be an Eclipse plug-in. For brevity, we omit names of 19 tools each of which was used in two papers.

One out of four papers introduces a new tool; another one out of four papers uses some home-grown tooling. Almost three out of four papers use existing tools.

The most popular standard tool, Eclipse—an IDE and a platform for plug-in development—is used in one out of seven papers. Other popular tools cater for source code analysis, clone detection, evolution analysis, data mining, statistics, quality analysis, document classification.

#### Home-grown tooling

ICPC	ESEM	CSMR	ICSM	MSR	WCRE	SCAM
5%	17%	20%	33%	33%	33%	42%

#### Introduced tools

ESEM	MSR	ICPC	SCAM	CSMR	WCRE	ICSM
4%	11%	19%	21%	30%	37%	44%

### E. Structural signs of rigorousness/quality

We do not aim to assess the quality or rigorousness of the studies. We capture presence of some of the aspects that are taken into account when assessing rigorousness/quality of research (cf., [23])—in that, we restrict ourselves only to the structural aspects.

1) *Study presentation aspects*: A clear set of definitions for the terms used in the paper is found in 25 papers. Research questions are adopted in 83 papers. In 22 papers, a “Goal-Question-Metric” approach is used. Explicit mention of null hypothesis or hypotheses is found in 23 papers. Section “Threats to validity” is present in 111 papers; of them, 75 discuss threats using classification described, e.g., in [24]: threats to external (mentioned in 73 papers), internal (59 papers), construct (53 papers), and conclusion (26 papers) validity.

If to consider combinations of these signs (definitions, research questions, hypotheses, and threats), the most popular one is the absence of all of them: demonstrated by 42 papers. The second most popular combination is presence of research questions and threats to validity: found in 34 papers. The third most popular—usage of only threats to validity—found in 29 papers. Together, these three combinations describe 60 % of the papers.

Half of the papers use research questions to structure their study. One out of seven papers uses a “Goal-Question-Metric” approach and/or formulate (null) hypotheses to structure their research. One out of seven papers provides an explicit set of definitions of the terms used in the study. Threats to validity are discussed in three out of five papers.

The following three combinations of structural signs describe at least half of the papers in each conference, except for WCRE, where only 44 % of papers are covered by these combinations.

No structural signs						
ICPC	MSR	WCRE	ESEM	ICSM	CSMR	SCAM
14 %	17 %	19 %	21 %	22 %	27 %	53 %

Both research questions and threats to validity						
ICSM	WCRE	SCAM	CSMR	MSR	ICPC	ESEM
8 %	11 %	16 %	20 %	22 %	24 %	42 %

Only threats to validity						
ESEM	MSR	SCAM	WCRE	CSMR	ICPC	ICSM
4 %	11 %	11 %	15 %	17 %	19 %	31 %

2) *Validation*: We captured the mentions of performed validation of done research. We have found evidence of some kind of validation in 88 papers. In 50 cases, validation was manually performed: either the results are small enough, or a sufficient subset is checked. In 27 cases, validation was done against existing or prepared results: actual data (when evaluating predictions), data from previous work, or an oracle/gold standard. In 8 cases, cross-validation was used.

### F. Reproducibility

We looked for signs of additionally provided data for a replication of the study. Since it is usually done via the Internet, we searched the papers for (the stems of) the following keywords: “available,” “download,” “upload,” “reproduce,” “replicate,” “host,” “URL,” “website,” “http,” “html”. In such manner, we have found links in 61 papers. In 6 cases, we could not find any mentioned material, tools or data,—links led to a general page or to a homepage, which we searched thoroughly but without success. In 3 more cases, we have found replication material on the website after some searching.

One out of three papers additionally provides online some data from the study, though not always to be found.

Additional data provided						
SCAM	ICSM	CSMR	MSR	WCRE	ESEM	ICPC
26 %	31 %	33 %	33 %	33 %	38 %	48 %

As to the nature of the provided data, in 25 cases, an introduced tool or tooling used in the research is provided. In 15 cases, the used corpus—in full or partially—is provided; the complete description of the corpus (list of used projects with their versions and/or links) is provided by 6 papers. Raw data is available for 14 papers; the same number of papers provide final or/and additional results of the study.

When the corpus is not provided by the paper, but the names of the used projects are mentioned, the main aspect of being able to reproduce the corpus is knowing which versions of the projects were used. We noticed that in 21 papers versions of the used projects are not provided. In 67 papers, versions of the projects are mentioned explicitly; in 26 more cases, it is possible to reconstruct the version from the mentioned time periods that the study spans.

Altogether, we judged 29 papers to be reproducible, meaning that either all components were provided by the authors or we concluded that the paper contains enough details to collect exactly the same corpus and the same tools. We did *not* judge if it is possible to follow the provided instructions, specific to the reported research.

We also would like to note that 8 papers mention that they are doing a replication in their study, of them 3 papers with self-replication.

We judged one out of six papers to be reproducible with respect to the used corpus and tools. We did *not* assess whether enough details were provided to re-conduct the research itself.

Judged to be reproducible						
ICSM	WCRE	SCAM	ICPC	ESEM	CSMR	MSR
3 %	4 %	16 %	19 %	25 %	27 %	33 %

### G. Assessment

Though usually information we extracted from the papers was scattered across different sections, half of the papers had



tables (listing projects, their names, versions, used releases, and similar information) that helped us during coding. We captured our confidence in the coded profile of each paper. For that, we used the following scale: high, moderate, and low level of confidence. The results are as follows: high—81 papers, moderate—78 papers, low—16 papers.

We have low confidence in one out of eleven papers that we have coded. In the rest, half of the time we are moderately confident and half of the time—highly confident in the results.

High confidence						
WCRE	SCAM	ICSM	ICPC	ESEM	CSMR	MSR
15 %	26 %	42 %	52 %	54 %	60 %	78 %

  

Moderate confidence						
MSR	CSMR	ESEM	ICPC	ICSM	SCAM	WCRE
17 %	33 %	33 %	43 %	53 %	58 %	67 %

  

Low confidence						
ICPC	ICSM	MSR	CSMR	ESEM	SCAM	WCRE
0 %	6 %	6 %	7 %	13 %	16 %	19 %

#### IV. RELATED WORK

We summarized related work—in the sense of other literature surveys on SE research—in the introduction and Table I. Thus, the key differences between our survey and previous work are these: i) predefined schema in previous work versus emerged schema in the present survey; ii) focus on specific forms or characteristics of SE research in previous work versus broad analysis of empirical evidence in the present survey. Below we compare the findings of related work where they overlap with ours.

As a general remark, we believe that our findings quantitatively differ from previous findings because of several factors: i) the dependence on the choice of venues: even conferences in our study differ considerably; ii) passed time: there is at least a five-year gap, during which popularity of empirical research and of its particular forms might have grown; iii) the cited papers use mostly journals: this may increase the aforementioned gap because of the longer process for journal publications; iv) snapshot versus longitudinal approach: we take into account all papers of the latest proceedings while the cited papers focus on a sample across several years.

The closest work to ours is by Zannier et al. [7]: they measured quantity and quality of empirical evaluation in ICSE papers over the years. Our work provides a snapshot study aiming to represent SE research broadly across conferences. Zannier et al. when assigning types to the papers, could confirm the self-classification of half of the studies. Which agrees with our observation that self-classification is rather weak among SE papers. They also observe the extremely low usage of hypotheses (only one paper) and absence of replications. We do find some adoption of null hypotheses and replications.

According to their classification, Glass et al. [5] have found 1.1 % papers to contain literature reviews and 3 % papers

to present “laboratory experiment (human subjects).” We also discover that number of literature surveys and experiments is low, but relatively it increased 2-3 times.

Kitchenham et al. [8] considered only 0.75% of surveyed papers to be systematic literature reviews. We have found literature surveys in 5 papers, one of which did not contain a clear methodology—a requirement to be met by Kitchenham’s inclusion criteria—leaving 4 papers. Thus, our percentage of detected literature surveys is 2.3 %

According to Sjøberg et al.’s study [6], only 2 % of the papers contain experiments, while we discover 10 % surveyed papers to contain an experiment. On the average, Sjøberg et al. detected an experiment to involve 30 participants—in 72.6% cases only students, in 18.6% cases only professionals, and in 8 % cases mixed groups. We have found that on the average an experiment involves 16 participants—in 57 % cases only students, in 14 % cases only professionals, and in 29 % cases mixed groups.

#### V. THREATS TO VALIDITY

1) *Choice of the papers:* We did not use journal articles—while they might provide more information or be of higher quality, we wanted to capture the state of the common research, of which we believe conference proceedings to be more representative. We have chosen conferences with proceedings of similar and reasonable size: so that not to skew the general results by one larger conference and so that to include all the papers but still be able to process them within reasonable period of time. Specifically, we excluded the ICSE conference, which had 87 long papers in the proceedings of 2012 edition. Altogether, this means our results might not be generalizable, but we believe them to be representative enough.

2) *Choice of the period:* Since we perform a snapshot study, it might be that some of the discovered numbers are a coincidental spike. A longitudinal study—possible future work—may provide more details and deeper understanding.

3) *Coding:* The effort was manual with occasional search by specific keywords (mentioned in the appropriate subsections of Section III). In 5 cases, papers were OCR-scanned.

**Human factor.** Coding was done by one researcher, but the results of the first pass were cross-validated during the second pass as well as during the aggregation phase. When in doubt, the researcher constantly referred back to the surveyed papers to double-check.

**Scheme.** We do not claim our coding scheme to be complete or advanced. We captured basic data related to the used empirical evidence, often either obvious or structurally supported. Therefore, we might miss sophisticated or under-specified forms of empirical research.

#### VI. CONCLUSION

In this paper, we presented a literature survey on empirical evidence in Software Engineering research.

*Answers to the research questions:* Coming back to the initial questions that motivated our research (see Section II), we suggest the following answers:

- I The overwhelming majority of surveyed conference papers use corpora—collections of empirical evidence.
- II The majority of the corpora consist of projects and can be characterized by size, code form, software language, evolution measures, requirements, and applied tunings.
- III There are no frequently used projects or corpora across all the papers. We have detected though some pattern of project recurrence with low frequency.

In what follows, we further interpret these findings.

*No “holy grail”:* Though corpora are used in the majority of the surveyed papers and some clusters of characteristics of the used corpora are recurrent (e.g., the use of many open source Java projects), the usage of established datasets is low. We suggest two possible reasons. First, adoption may be low only yet: among detected datasets being used (see Table IV), the oldest dataset, SIR, was introduced in 2005, the youngest, Qualitas—in 2010. Second, researchers may prefer to collect and prepare their corpora themselves, because there might not be a “holy grail” among corpora to suit all possible needs. Partially, this assumption is supported by the fact that, even on the level of projects, no clear favorite was detected among the papers. The emerged schema with its components for requirements and tunings for corpora also substantiates indeed the different needs of research efforts.

*Community-specific curated collections:* On the other hand, we find that three out of seven conferences have favorite projects, when considered separately—projects that are used by a quarter of the papers within these conferences. This leads to a refined version of the third question in our study: When it is possible to detect commonly used projects within a conference, would it be useful to provide a curated version of them? Generally, it is clear that even requirements and tunings are recurrent across research efforts, and hence, some “product line” of curated collections and some discipline of “corpus engineering” may ultimately lead to more reuse of empirical evidence. These are topics for future work.

*Top-down vs. bottom-up introduction of methodology:* While there is a need for adoption of advanced and theoretically specified forms of empirical research, we believe that there is a certain amount of de facto empirical research in Software Engineering that has formed historically. This survey sought to understand the characteristics of empirical evidence in research—also to enable assessment, if not improvement, of research quality. Future work includes aligning research areas or goals with the kind of used empirical evidence: deeper understanding of the needs may provide insights for streamlining research.

## REFERENCES

- [1] R. Lämmel, S. Kitsis, and D. Remy, “Analysis of XML schema usage,” in *Conference Proceedings XML 2005*, 2005.
- [2] R. Lämmel, E. Pek, and J. Starek, “Large-scale, AST-based API-usage analysis of open-source Java projects,” in *SAC*, 2011, pp. 1317–1324.
- [3] R. Lämmel, R. Linke, E. Pek, and A. Varanovich, “A framework profile of .net,” in *WCRE*, 2011, pp. 141–150.
- [4] R. Lämmel and E. Pek, “Understanding privacy policies - A study in empirical analysis of language usage,” *Empir. Softw. Eng.*, vol. 18, no. 2, pp. 310–374, 2013.
- [5] R. L. Glass, I. Vessey, and V. Ramesh, “Research in software engineering: An analysis of the literature,” *Inform. Software Tech.*, vol. 44, no. 8, pp. 491–506, 2002.
- [6] D. I. K. Sjøberg, J. E. Hannay, O. Hansen, V. By Kampenes, A. Karahasanović, N.-K. Liborg, and A. C. Rekdal, “A survey of controlled experiments in software engineering,” *IEEE Trans. Softw. Eng.*, vol. 31, no. 9, pp. 733–753, 2005.
- [7] C. Zannier, G. Melnik, and F. Maurer, “On the success of empirical studies in the international conference on software engineering,” in *ICSE*, 2006, pp. 341–350.
- [8] B. Kitchenham, O. Pearl Brereton, D. Budgen, M. Turner, J. Bailey, and S. Linkman, “Systematic literature reviews in software engineering - A systematic literature review,” *Inf. Softw. Technol.*, vol. 51, no. 1, pp. 7–15, 2009.
- [9] B. G. Glaser, *Doing Quantitative Grounded Theory*. Sociology Press, 2008.
- [10] S. Adolph, W. Hall, and P. Kruchten, “A methodological leg to stand on: Lessons learned using grounded theory to study software development,” in *CASCON*, 2008, pp. 13:166–13:178.
- [11] D. I. K. Sjøberg, T. Dybå, and M. Jørgensen, “The future of empirical methods in software engineering research,” in *FOSE*, 2007, pp. 358–378.
- [12] M.-A. S. Steve Easterbrook, Janice Singer and D. Damian, “Selecting empirical methods for software engineering research,” in *Guide to Advanced Empirical Software Engineering*, F. Shull, J. Singer, and D. I. Sjøberg, Eds. Springer, 2008, pp. 285–311.
- [13] J.-M. Favre, D. Gašević, R. Lämmel, and E. Pek, “Empirical language analysis in software linguistics,” in *SLE*, 2011, pp. 316–326.
- [14] H. Do, S. Elbaum, and G. Rothermel, “Supporting controlled experimentation with testing techniques: An infrastructure and its potential impact,” *Empir. Softw. Eng.*, vol. 10, no. 4, pp. 405–435, 2005, <http://sir.unl.edu/>.
- [15] “MSR mining challenge 2008, 2011,” <http://msr.uwaterloo.ca/msr2008/challenge/>, <http://2011.msrfconf.org/msr-challenge.html>.
- [16] Y.-G. Guéhéneuc, “PMAR: Pattern-like micro architecture repository,” in *EuroPLoP Focus Group on Pattern Repositories*, 2007, <http://www.ptidej.net/download/pmar/>.
- [17] “The PROMISE repository of empirical software engineering data,” since 2005. <http://promisedata.googlecode.com>.
- [18] E. Tempero, C. Anslow, J. Dietrich, T. Han, J. Li, M. Lumpe, H. Melton, and J. Noble, “Qualitas corpus: A curated collection of Java code for empirical studies,” in *APSEC*, 2010, pp. 336–345, <http://qualitascorpus.com/>.
- [19] M. D’Ambros, M. Lanza, and R. Robbes, “An extensive comparison of bug prediction approaches,” in *MSR*, 2010, pp. 31–41, <http://bug.inf.usi.ch/>.
- [20] S. M. Blackburn, R. Garner, C. Hoffman, A. M. Khan, K. S. McKinley, R. Bentzur, A. Diwan, D. Feinberg, D. Frampton, S. Z. Guyer, M. Hirzel, A. Hosking, M. Jump, H. Lee, J. E. B. Moss, A. Phansalkar, D. Stefanović, T. VanDrunen, D. von Dincklage, and B. Wiedermann, “The DaCapo benchmarks: Java benchmarking development and analysis,” in *OOPSLA*, 2006, pp. 169–190, <http://dacapobench.org/>.
- [21] E. Linstead, S. Bajracharya, T. Ngo, P. Rigor, C. Lopes, and P. Baldi, “Sourcerer: Mining and searching Internet-scale software repositories,” *DMKD*, pp. 300–336, 2009, <http://sourcerer.ics.uci.edu/>.
- [22] B. Kitchenham, “Procedures for undertaking systematic reviews,” Keele University and National ICT Australia Ltd., Tech. Rep. TR/SE-0401, 2004.
- [23] M. Ivarsson and T. Gorschek, “A method for evaluating rigor and industrial relevance of technology evaluations,” *Empir. Softw. Eng.*, vol. 16, no. 3, pp. 365–395, 2011.
- [24] C. Wohlin, P. Runeson, M. Höst, M. C. Ohlsson, B. Regnell, and A. Wesslén, *Experimentation in software engineering: An introduction*. Kluwer, 2000.